



Altering Theories of Learning and Action: An Interview with Chris Argyris

Author(s): Chris Argyris and Mary Crossan

Source: *The Academy of Management Executive* (1993-2005), Vol. 17, No. 2 (May, 2003), pp. 40-46

Published by: [Academy of Management](#)

Stable URL: <http://www.jstor.org/stable/4165953>

Accessed: 12/06/2014 13:52

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Academy of Management is collaborating with JSTOR to digitize, preserve and extend access to *The Academy of Management Executive* (1993-2005).

<http://www.jstor.org>

Altering theories of learning and action: An interview with Chris Argyris

Interview by Mary Crossan

Executive Overview

Chris Argyris received his Ph.D. in organizational behavior from Cornell University in 1951 and served on the Yale University faculty for the next twenty years. He then became the James Bryant Conant Professor of Education and Organizational Behavior at Harvard University and held joint appointments over time at the Business, Law, and Kennedy Schools. His early research dealt with organizational structures and organizational change, in such works as *Integrating the Individual and the Organization* (1964) and *Organization and Innovation* (1965). His primary research methodology has been intervention, and his ongoing concern has been the applicability of knowledge. His work on individual and organizational learning culminated in the classic book which is the subject of our Retrospective.

Argyris is the author of 33 books and monographs and over 400 articles. He has been awarded eleven honorary doctorates. Other honors include the Academy of Management's Irwin Award, the Kurt Lewin Award, the American Psychological Association's Gold Medal Award for Life Achievement in the Application of Psychology, and several other "lifetime achievement" awards. He has been elected a Fellow of the Academy of Management, the American Psychological Association, the National Academy of Human Resources, the Canadian School of Management, and the International Academy of Management. The Chris Argyris Chair in Psychology has been established in his honor at Yale University.

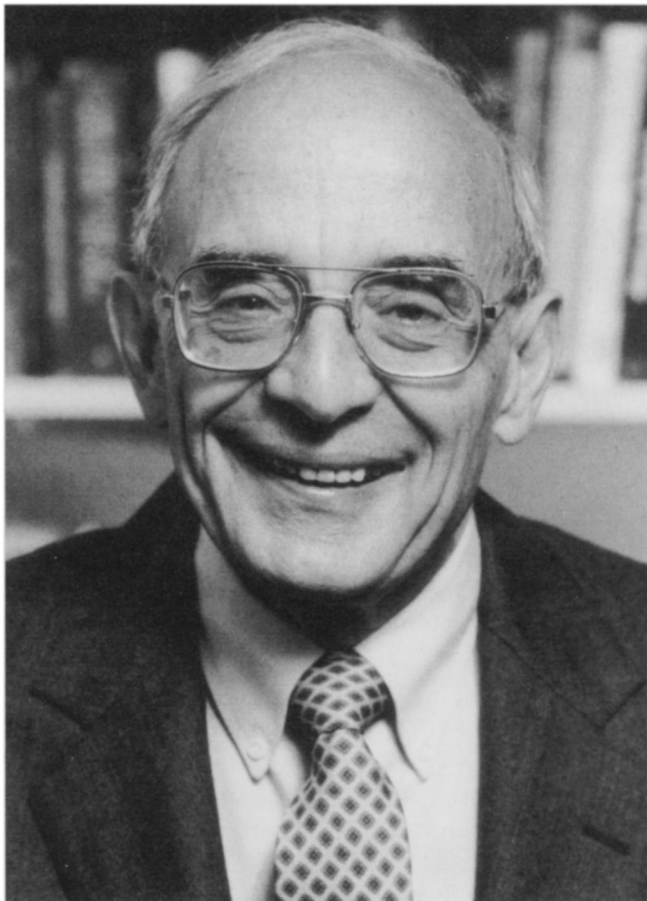
Perhaps you could start by taking us back to 1978 and setting the context for your book with Donald Schön, *Organizational Learning*. You noted in one interview that publishers had said, "We know you and we know Don, and we respect both of you, but do you think this topic will ever be of interest to the business community?"

Don and I began collaborating in 1974, and we became increasingly focused on learning. Our theory was on learning as the detection and correction of error. Our research moved from individual to group, and by the time this book was being written, we were telling ourselves that we ought to write something that includes organizational learning. I must say that the quote is accurate, but we never worried about it. We believed that organizational learning would become increasingly more important, and that double-loop learning would be the bigger challenge.

There had been so much work in behavioral and cognitive psychology on individual learning. Did you feel that organizational learning was fundamentally different from individual learning in an organizational context? As you moved to the organization level, what did you draw upon?

I think that it is important to distinguish between *enabling* organizational learning and *producing* it. Enabling organizational learning includes group, inter-group, and organizational features such as policies, practices, rules, and organizational memory. Producing organizational learning is done by individuals taking action. Whenever individuals produce organizational learning, they act as agents for the organization.

Our shift was to show that some of the better work in individual learning could in fact move us toward organizational learning. When you ask "How do you know when you know something?"



Chris Argyris

the answer is "Can you produce what you claim you know?" So the operational criterion for learning was detection and correction of error. But you would not know that you had corrected it without implementation. In social psychology there is the notion of internal validity and external validity. With internal validity you design your research to remove any confounding variables that would hurt your generalization. And external validity claims that learning in a research context is applicable to other contexts; it has external relevance. We began to think that this perspective limits science because it suggests there is no reason to focus on implementable validity. People studying cognition, because they were making claims about how the mind-brain works, almost always gravitated to questions of implementable validity. Many did research that produced what they were talking about, so I would say that I was—and still am—importantly influenced by cognitive theory.

So, if the cognitive dimension is more akin to, or more open to being tested for, implementable validity, it then becomes more practical. But now I add something else. The implementable validity should be producible by a scholar or by a practi-

tioner—it must be applicable to everyday life conditions and, again, both the scholar and the practitioner should be able to use it.

Early in my life I wrote a book called *Personality and Organization* in which I identified a number of problems. Then I asked myself, "Well, how do we change this? And how do we correct it?" And these problems ranged from the enrichment of jobs to the redesign of control systems such as accounting procedures. But as I looked at the problems, questions arose about who had control over these activities. I interviewed managers and discovered they were unaware—it was tacit.

So I wanted to get at the way management reasoned. My approach was to begin at the top and go no further until there was evidence that their defensive reasoning mindsets had changed to more productive reasoning mindsets. Now, how do you do that? Do you put every executive through therapy? That would not work. At that time the T-groups I was involved with were still more clinical than cognitive. The next question was, "Can we create groups where people could learn from each other, especially if they were an organic team?" In the early days I thought that groups had to be less than 20 people. The further I moved toward cognitive theory, the bigger the groups could become. In order to study 100 people or more, I had to build a model that claimed to know what goes on in the human mind. How do you produce behavior? Through the use of your mind. How does the mind produce behavior? By designs that are programmed into the mind, which can be warehoused and then retrieved. What are the designs? They are causal designs—doing something causes something else to occur. Also, the mind needs evidence that what it claimed it could do actually happened. What emerged was not only a causal theory of how to do this but a causal theory of how to go *meta* and study how well the process worked. The clearer we can be about that, the more effectively we can help people learn—help get executives to learn and change.

I wanted to get at the way management reasoned.

Human beings are skillful at two kinds of reasoning. One I call productive reasoning and the other I call defensive reasoning. Now, the function of productive reasoning is to seek, as best we can, truth about the effectiveness of what we do and what we claim. This gives rise to causality, testing, the notion of transparency, testing methodology, and so on. Defensive reasoning, serving as self-protection, leads to self-referential logic.

And people are skillful at both. They learn productive reasoning when they study various kinds of managerial disciplines such as accounting, finance, and economics. They learn defensive reasoning, ironically, in the conduct of social science research. The fundamental assumption behind good social science is that it describes, as accurately as possible, the universe that researchers have selected to study. So its function is descriptive. We teach less about its normative function, even though the world we are studying is all normative. The people we study and describe behave normatively.

I believe that it is not only important to describe what we observe, but to describe what the universe would do if someone really wanted to change it. This is a description we don't focus on because nobody tinkers with it in the double-loop. Therefore social science, whether unintentionally or tacitly intentionally, aims to describe the universe as it is, but misses the portion that is underground—not in our focal point because we are not working on changing it. These ideas led me to questions about intervention, double-loop learning, and so on.

Don and I also claimed that in each human being there exist theories of action and that these, too, are normative. There is no reason to believe that one set of governing values is more objective than the other. We said there is a difference between the theory that people espouse and the one they actually use. We framed this and called it Model I. I then tried to show that for rigorous research the theory-in-use (not the espoused theory) is consistent with Model I.

How can we help researchers see that their research exemplifies a theory that they deny? I took a look at some of the best research—Carl Hovland and his group years ago had studied communication and trust. How do you communicate in such a way that you are trusted? In effect, one of the implications for action was if you are communicating with a group that you believe is bright, give them two or more alternatives, but not too many. They do not trust a one answer. If you think that the group is not so bright, give them one answer. And I remember asking Carl, "Would you get up in front of a group and say 'Since I believe you're dumb, I'm going to give you one answer'?" He said, "Of course not." I said, "So you're covering up?" and he responded, "Well, of course." So I asked, "Who teaches you to cover up, and who teaches you to cover up that you're covering up?" He indicated he would do so "in the name of truth," becoming a Nixonian in the name of science. He said he had not thought of that. People tend to provide answers that are not what we would expect from rational people. For example, their answers often have to

do with the recency of experience. Recency is a heuristic they use. The interesting thing about those heuristics is that they are valid in a Model I world. Then the question is, "Why don't social scientists think about creating another world?" If they ever do, they will not only have to describe the world as it is, but also a world that might be. It is doubtful that the heuristics identified in research would operate in a Model II world.

How can we help researchers see that their research exemplifies a theory that they deny?

Do you think that is because social scientists are not trained and, perhaps, not motivated to even think about what an alternative world might look like? Are they trained as researchers in the Model I tradition?

People who are attracted to social science and who act consistently with Model I theories in use tend to feel safe because they are in unilateral control. I believe they are frightened to do interventions that don't permit them that kind of cover. In an intervention you can claim something and the other person can disagree. In an experiment, subjects rarely tell the researcher that he is wrong. I remember one world-class social psychologist saying, "Chris, I don't like intervention because it scares me to hell." Here was a world-class scholar with loads of positive aggressiveness and strength. I asked him why he was frightened about this. And he said, in effect, that he'd never thought of this. But he added, "If I had thought that social science meant putting myself at risk in front of the people I call subjects, I'm not sure I would have done social psychology or sociology. I might have done neurobiology where there's less risk. The neutrons don't confront you or create embarrassing situations." So, scholars don't think about intervention, and they don't want to take the associated risk. And the people we attract to the field are people who feel comfortable about not thinking about it, people who, when asked why they don't think about it, defend themselves by saying "What, are you kidding? That's the methodology."

In the Introduction of the book you state that:

Our ultimate goal is to help individuals unfreeze and alter their theories of action so that they, acting as agents of the organization, will be able to unfreeze the organizational learn-

ing systems that also inhibit double-loop learning. In defining this goal, we realize that we are probably biting off much more than we can possibly chew. Not only do we have the temerity to question underlying human theories of action and organizational learning systems, but we are calling into question some of the most basic societal norms and values. Moreover, we even strive to present new models of action for individuals, organizations, and societies.

One reason for taking the risk that some readers may view our goals as premature or naive, is that we believe that social scientists should explore possible resolutions of these learning issues. Unless people acting as agents for organizations and societies are able to learn how to detect and correct double-loop errors, the survival of the society may be in doubt (p. 5).

Is one of your underlying motivations to think about the grand picture of this society?

Yes. I have two types of grand picture. If our position is correct, there is a kind of built-in entropy that human beings create for their world. This entropy is caused by Model I defensive routines that feed back, creating a syndrome that is circular, self-sealing, and non-testable. I see the management of the world as slowly deteriorating. I was interested in what might lead to a healthier life, a healthier, more effective organization—be it for-profit or non-profit. I did have a notion that we should put social science to the task of keeping the world from going to hell by its own behavior.

I did have a notion that we should put social science to the task of keeping the world from going to hell by its own behavior.

I had another, related interest. What I loved about doing work in social science was seeking some kind of truth—truth with a small *t*. But the more I studied it and its methodology, the more I became interested in justice. Was there justice in putting people through experiments for the sake of learning, no matter what the subject or its aims? I came to the notion that someday there will be a point where I cannot say truth is good no matter what because it may violate the issue of justice. I have not fully reconciled these principles and, in my opinion, this is another reason why social scientists ought to learn about the importance of justice. They should

ask themselves whether their research keeps one eye on the validity of their conclusions and the other eye on the implications for justice.

In your book you state, “There has probably never been a time in our history when members, managers, and students of organizations were so united on the importance of organizational learning.” And you go on to talk about those conditions that merit the need for organizational learning. Does this context still apply today? Is there progress? Are we moving in the right direction as a society?

It is a question I keep asking myself. I want to go back to the concepts of productive reasoning and defensive reasoning, and I want to say that the world is becoming more productive in its reasoning and more defensive—it is the same people who are moving in both directions. It is not that one group is more productive and the other is more defensive. I believe that we are making progress in the use of productive reasoning. I also believe that we are increasing our use of defensive reasoning, especially where the correct use of productive reasoning would create conflict and dilemmas. We use defensive reasoning under conditions when it is least effective in solving the problems at hand.

I now want to back up and make one correction. There are people who are educated in, or stronger at, productive reasoning. Looking at the development of accounting, finance, marketing, all the managerial disciplines, and some features of HR such as compensation schemes and so on—these are mostly based on productive reasoning. I used to observe people teaching accounting thirty years ago. Accounting is now much more rigorous. You can’t get away with defensive reasoning today as much as you could years ago. Years ago, in an organization, being taller, better looking, and male made a difference. One reason these things mattered was the fact that the discipline was not particularly rigorous.

The information science revolution has accelerated the whole process. Now the fundamental notion of information science is that truth is a worthy goal and how we arrive at it ought to be transparent. The techies say “garbage in, garbage out.” However, as human beings, we embrace truth only when it is not threatening or embarrassing. Otherwise we hide it, or massage it. IT creates a world in which what goes on underground, and is defensive and not transparent, can be made more transparent. The audience can ask, “Wait a minute, where is this going?” IT will help to counter the art of spin. Television has done a good job of turning people

off. After a while they say, "Oh no, not again." I predict that the IT world will become more powerful. People will become better at recognizing spinning, partially because they have become inundated and partially because they are spinners themselves. I predict that we will see an increasing degree of conflict and that it will be less possible to cover up than it was 30 years ago. The new skill will be spinning. But that, too, will have limits. People will get tired of it. I predict that productive reasoning will become more powerful—not because of churches, schools, and so on, but because of a cognitive invention called information science which has the great strength of being based on how the mind works.

A manager showed me a computer screen of what happened the prior day in 52 plants throughout the world. And he said, "Now all this information is public to all people at all levels of management so they have the right to use it too, and if they don't use it correctly and I find out that they don't, I'm going to get after them. And if I don't use it correctly and they find out, they should get after me." So you see how the computer is creating a world that is much more information friendly than the world that used to be.

It seems to me that there is no guarantee, even with information systems, that we will have valid information. The information may be more of the single-loop kind—the information collected around preserving the existing order. I would place more hope on developing the basic capacity to reason and think, through something like the study of philosophy.

I agree that the examination of philosophy is important. I am a devoted student of the philosophy of science, and Don had a Ph.D. in philosophy. And I agree that IT is used primarily to deal with single-loop issues. IT may make it more likely that managers and scholars can be more rigorous in understanding and explaining, but it is less helpful in producing new prescriptions that require a double-loop perspective. No matter how good a computer program is, it is not guaranteed to be error-free. And if there is a guarantee, it is probably more for the single-loop routine stuff. One way to respond is to get comfortable with going *meta*. Learn to reflect on whether you are unknowingly kidding yourself, despite the printouts. Ask yourself to what extent you unknowingly maintain the status quo when you deny being self-deceived.

What does it mean to go *meta*? You take a look at what you have just done—you reflect on it. Be a

"reflective practitioner," as Don nicely called it. The reflective practitioner is good at going *meta*. But what does it mean to be good at going *meta*? This is something that philosophy, it seems to me, can teach people—teach them to think about governing values and help them think about criteria for how to test a claim so that there is no error. For me, philosophy would be the key feature in vigilant testing of IT programs.

Can we pursue the set of conditions of error and corrective responses that you noted in your book? How do those conditions arise? Take "information withheld," for example. Why do some people not present their reasons, and why don't others ask?

First, the present notion of effectiveness, often, is to be good at Model I. Under Model I we win, don't lose, don't encourage inquiry, and so on. Second, such reasoning creates an organizational defensive routine that feeds back to reinforce the Model I behaviors. People withhold information—they massage it—because they see that as a sign of effective leadership. If they knew another kind of leadership, they would not have thought of it.

People withhold information—they massage it—because they see that as a sign of effective leadership.

We educate youngsters early in life about Model I. We educate them in defensive routines in adolescent groups, college groups, and in organizations. They become skillful at it—the business of withholding information, spinning, and so on. It is tacit, taken for granted and, indeed, they are rewarded for doing it. Our view is that we create conditions where error is not confronted appropriately. But the people who are doing it, they think it is quite appropriate. If you ask them "Aren't you doctoring this information?" they respond, "Well, Chris, that's true of all organizations." Yes, but under what conditions can we begin to change that? And that is where IT fits in. IT makes it possible to interrupt the routines—if the IT is the thoughtful kind that you are talking about, if it is employed and continually taken *meta*, and if you are vigilant about your own reasoning, then it seems to me that it is less likely you can hide from data that questions your conclusion.

So, what is it that we can do to create a better platform for Model II organizations if we have to undo a lot of things learned at an early age? We can certainly start at business schools, but what about earlier than that—what do we need to do?

The place to start is at school, when the children are young. When I came to Harvard, I had a dual appointment, but I committed myself to major work on schools. The teachers I had in school then—and the students—espoused a Model II world but became increasingly upset over the requirements of such a world. Interestingly enough, looking back on it, I think that one of the things that most frightened them was the concept of critical reasoning. In an ideal world, the study of critical reasoning would begin in schools, but after 18 years I decided that would take a long, long time. I decided I'd have to be more explicit about what it means to be a teacher of students—I would include what we are now talking about. And the teachers never got that kind of advice. The few that did were quite enamored with it; the intellectually stronger ones said "Great" and the intellectually weaker ones said, "Gee this is inhumane—you're being too scientific, Chris. Life is not that rigorous."

So, having a joint appointment and teaching in the business school, I found that the faculty was much more willing to look at my ideas, and their limits. Going meta on accounting and so on was okay—it was in the literature. What was not in the literature was going meta in front of the students. So I would say I would start at schools. As well, I would have executive courses in schools and in businesses that focused on the defensive mindset of executives. There is an article coming out in *Academy of Management Learning and Education* which illustrates how I do that.

Is there ever a time when error should not be confronted? When should we not question the fundamental rules we operate by? Is there any condition under which we would not act in a Model II way?

Yes—when there is objective data showing that the decision you must make requires time that you don't have. Someone quoted me saying, "It doesn't make any sense to die in the name of Model II." So there are conditions under which it would be productive to use Model I. But once you do it, you must then find time to reflect. In the U.S. Army they have created feedback sessions to reflect on what happened. I did not observe them directly, but I saw videotape recordings of these feedback sessions. People were quite open about errors that they, or

others, had made. The next question was, "How do we guard against those errors?" From talking with the officers, I learned that not only did the sessions reduce some of the ambiguity about what to do, but they also reduced the ambiguity about how fearful it can be to make these decisions. Participants said, "I don't know what to do, but I know the Major and he's not going to get teed off at me for doing this." So they had a sense of stewardship being built. However, their behavior was more consistent with Model I. It would be interesting to experiment to create a feedback session where the participants used Model II theory-in-use. Or they would say, "I do know he's going to be upset. I don't have the time, and he knows I don't have the time. He's going to judge me on how well I use my reasoning processes, and I'm willing to accept that." I believe that the use of Model II and productive reasoning give us the best chance to identify those conditions; to own up to them; to encourage challenging our views.

If you think of yourself as one of the parents of the genealogy of work on organizational learning, what are you most pleased about at this time? What do you see as the most important current debates? What are the main dangers—perhaps defensive routines—of our research in organizational learning?

I am pleased with the increasing recognition of the centrality of learning. I think much of the credit should go to IT. When I was growing up, managers used to say they *hired a hand*, and they really meant it. They did not mean they hired a *mind*. But today they do say they hire minds. In a world where minds are hired, learning becomes central. I used to teach newly commissioned admirals in the Navy and Marine Corps generals. The first time I tried, I thought, "What am I doing here talking about Model II?" Well, during the first class everything was going fine until, suddenly, one brigadier general raised his hand and said, "Professor, I don't know what the hell you're talking about. I'm going to be honest with you. If the President of the United States told me that all my men should shave their heads and go and take that hill, I can assure you that we'll all do that." The other brigadier general there said to him, "Bill, that's dead. If we're going to have the right people, we need to use the machinery and the minds that are required to use the machinery."

So the importance of learning pleases me. But what I am especially pleased about is the connection to effective action in ways that specify causality. Don and I agreed that the literature on organizational learning does not make this connection.

The much smaller literature on learning in organizations does—it is especially strong at the individual and group levels and is becoming stronger at the higher levels of complexity. The danger is in becoming information-happy—inundated with information—and using that as an excuse not to take action, or to take action by studying things. People may believe that being inundated with information gives them a valid reason to refrain, resist, or delay making a decision. And they can often get away with it by saying, “Yes, sir, but we don’t have the data. We need more data.” We need people who are not frightened by incompleteness and ambiguity, who will develop heuristics of the kind that will enable them to make decisions. To do this they will require skill at meta learning along with skill at decision-making.

The danger is in becoming information-happy—inundated with information—and using that as an excuse not to take action, or to take action by studying things.

When you think about further work that will enhance either the value of these approaches for practitioners in all sorts of organizations, or their value for more general social science theories of human action, what are you most hopeful for?

I believe that social scientists have the responsibility to conduct research that has external and internal validity *and* that places implementable validity at the same level of importance. This means that we should produce knowledge that is implementable, which in turn means that its propositions must meet the features of a design for action—it must specify the consequences and the behavior or means to reach it. Add to this a meta specification of how well you did. I mentioned earlier that implementable validity has not been a strength of social science, and it ought to be.

Returning to the earlier discussion about implementable validity as well as the notion of going beyond description to what might be possible—do these suggest a need for more research interventions?

Yes, I would agree with that. My view is that intervention is a very critical skill and, interestingly enough, we don’t teach our students very much about it. If we do, it is intervention at the single-

loop level, and not at the double-loop level. I went to Denmark with a group of IT scholars. They had a study on IT—it was a particular kind of test to be used for an evaluation of the effectiveness of IT in an organization. These people produced a paper that, among other things, showed beautifully the defensive routine of IT and line management and how it made IT less effective. It also showed management denying responsibility for this. The researchers admitted that management needed more double-loop learning, but they stopped there. One reason is that the faculty members didn’t know how to do double-loop intervention.

In your view what are the key contributions of *Organizational Learning*? Have scholars used the book in the way you expected? Are there messages that have been lost that you would like to emphasize, or others that have been overemphasized or poorly understood?

In the poorly understood category, I’d put the lack of research on double-loop learning. The largest amount of work that has been done is on single-loop or routine issues. I think that limits social science as a discipline and as a contributor to a better world. The notion that I think has been misunderstood is that somehow there is a difference between individual learning and organizational learning. I have read and I think I understood the critique. Think of insight and awareness—the new knowledge they create becomes an enabler for taking action. But if you use the criterion that we use—that not only do you get the insight, but you actually produce it—then I can’t see how anybody but individuals produce it. This brings us back to our original stance, so to speak. I can see how groups produce something. But if you think about Irving Janis’s work on groupthink, you can see how individuals acting with each other create the group property.

Chris, thank you for taking the time to share your insights with us. It has been a pleasure discussing the wide-ranging ideas and implications arising from this seminal book.

Mary Crossan holds the Donald K. Jackson Chair in Entrepreneurship and is an associate professor of strategic management at the Richard Ivey School of Business, University of Western Ontario. Her research on organizational learning, strategy and improvisation has been published in such journals as the *Academy of Management Review*, *Journal of Management Studies*, *Organization Science*, *Organizational Dynamics*, and the *Strategic Management Journal*. Contact: mcrossan@ivey.uwo.ca.